

## LA-UR-17-30493

Approved for public release; distribution is unlimited.

Title: ATOM - Accelerating therapeutics through opportunities in medicine

Author(s): McMahon, Benjamin Hamilton  
Dotson, Paul Jeffrey

Intended for: Presentation

Issued: 2017-11-16

---

**Disclaimer:**

Los Alamos National Laboratory, an affirmative action/equal opportunity employer, is operated by the Los Alamos National Security, LLC for the National Nuclear Security Administration of the U.S. Department of Energy under contract DE-AC52-06NA25396. By approving this article, the publisher recognizes that the U.S. Government retains nonexclusive, royalty-free license to publish or reproduce the published form of this contribution, or to allow others to do so, for U.S. Government purposes. Los Alamos National Laboratory requests that the publisher identify this article as work performed under the auspices of the U.S. Department of Energy. Los Alamos National Laboratory strongly supports academic freedom and a researcher's right to publish; as an institution, however, the Laboratory does not endorse the viewpoint of a publication or guarantee its technical correctness.

# ATOM – Accelerating Therapeutics through Opportunities in Medicine

**LANL capabilities in high performance  
computing and biological sciences**




**Paul Dotson**  
Deputy Associate Director  
Theory, Simulation, and  
Computation

**Ben McMahon**  
Deputy Group Leader  
Theoretical Biology

November 17, 2017

# LANL Organizations

## Bioscience spans 9 Divisions across LANL





# LANL: Leading Edge of Computing in Service to the Mission



CDC 6600  
1966

[Small/large  
core memory]



IBM Stretch  
begat IBM  
360 1962



MANIAC I  
1952



IBM 405  
1943



Cray 1  
1976

[Vector  
machine]



Cray XMP  
1983

[SMP]



TMC CM5  
1992

[MPP/Data  
Parallel]



SGI Blue  
Mountain  
1998

[Massively process  
parallel]



LNXI  
Lightning  
2014

[commodity computing]



Roadrunner  
2007

[Hybrid architecture]



Trinity 2016

[Massive Memory \  
Burst Buffers]




Dwave Ising  
2016

[Quantum  
Annealing]

- LANL's HPC mission has always been two-fold:
  - Provide HPC to execute scientific mission
  - Advance HPC to help execute future scientific mission
- Core NW mission needs have been major industry driver, but that has changed.
- Significant changes in architecture have accompanied the increasing power
- Resulted in rich capability of coupling scientific algorithms to varied architectures (i.e. scaling, messaging, and vectorization)

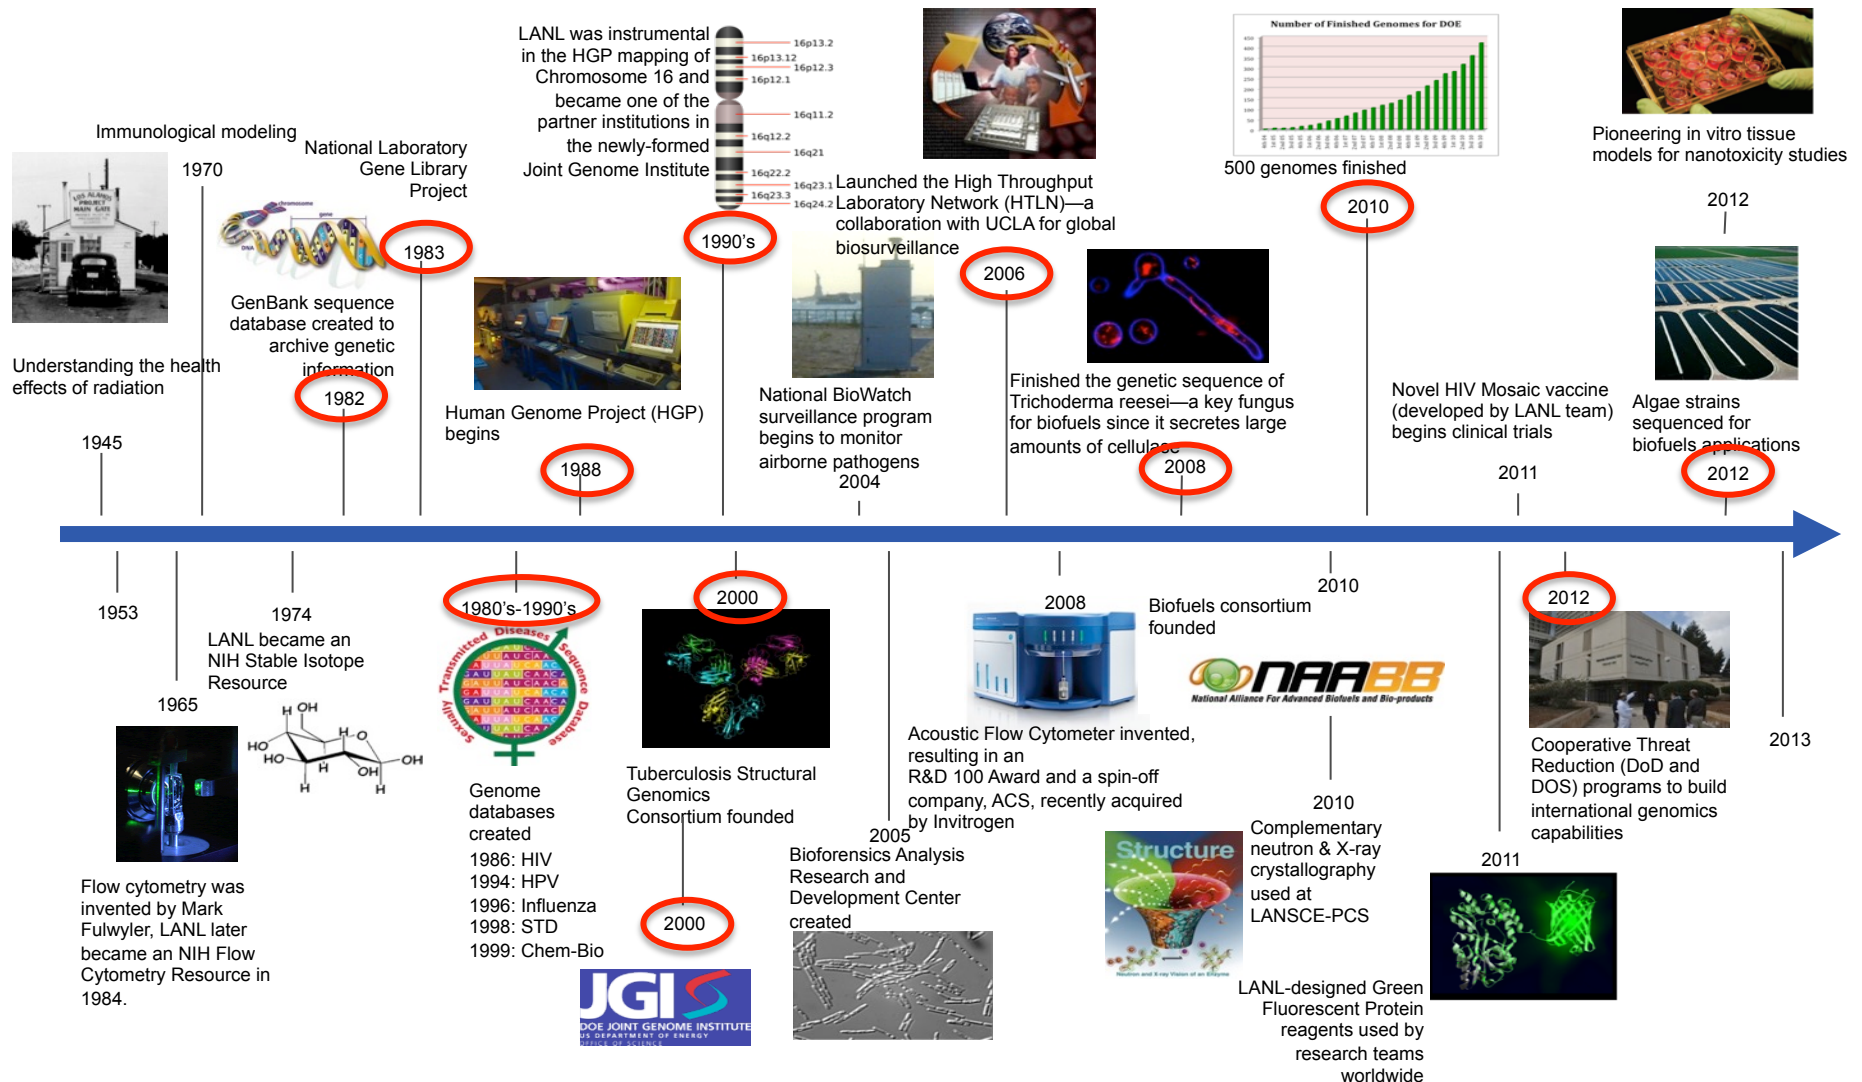
# Computing Has Always Been a Core Component of the Weapons Programs; the Landscape Changed in 1992

Number of U.S. Nuclear Events



Processing Power

# A timeline of LANL contributions to biological science





## ATOM: Accelerating Therapeutics for Opportunities in Medicine

Create a new paradigm of drug discovery that would reduce the time from an identified drug target to clinical candidate from the current ~6 years to just 12 months.


ATOM will develop, test, and validate a multidisciplinary approach to drug discovery in which modern science, technology and engineering, supercomputing, simulations, data science, and artificial intelligence are highly integrated into a single drug-discovery platform that can ultimately be shared with the drug development community at-large.



# From RCT to retrospective study. A modeling requirement is born

RCT have been the gold standard for demonstrating effectiveness, but is increasingly untenable as the situation becomes more complex (higher dimensionality)

But even RCTs can suffer from selection bias and design flaws, while, with proper modeling, retrospective analyses not only are accurate, but also reflect real-world settings.




**Figure 1.** Range of Point Estimates According to Type of Research Design.

The studies evaluated bacille Calmette-Guérin vaccine and active tuberculosis (13 randomized, controlled trials and 10 case-control studies), screening mammography and mortality from breast cancer (8 randomized, controlled trials and 4 case-control studies), cholesterol levels and death due to trauma among men (4 of 6 randomized, controlled trials [2 trials did not provide point estimates]; the results of the 14 cohort studies were not reported individually), treatment of hypertension and stroke among only the men in the studies (11 randomized, controlled trials and 7 cohort studies), and treatment of hypertension and coronary heart disease among only the men in the studies (13 of 14 randomized, controlled trials [1 trial did not provide point estimates] and 9 cohort studies). Solid circles indicate randomized, controlled trials, and open circles observational studies.

**RANDOMIZED, CONTROLLED TRIALS, OBSERVATIONAL STUDIES,  
AND THE HIERARCHY OF RESEARCH DESIGNS**

JOHN CONCATO, M.D., M.P.H., NIRAV SHAH, M.D., M.P.H., AND RALPH I. HORWITZ, M.D.



# Overview: Uncertainty quantification through integration of diverse information

## Physical Models

Bacterial transcription initiation  
Gene silencing on histones  
Cell signaling pathways  
Membrane transport  
Alzheimer's disease & MD

## Data Analytics

HIV databases  
Mosaic vaccines  
Cancer genomics  
Signature-based  
metagenomics

## Decision Support

HCV therapies  
Pandemic simulations  
DoD Biosurveillance  
VA Patient trajectories

The above examples of past work illustrate how models improve understanding of biological problems related to pharmaceutical development

Athena: Experimental multi-organ system for toxicology

High Performance Computing: Practical foundation rooted in multi-sponsor experience

Athena and HPC describe empirical background and HPC capability that will be needed.



# Physical models


- Bacterial transcription initiation complex
- Human gene silencing of DNA on histones
- Physically-based models of cell signaling
- Membranes and transport of Gram negative bacteria
- Relating simulations of amyloids to Alzheimer's disease

# Molecular modeling of protein complexes enables interpretation of bioinformatics

The bacterial RNA polymerase complexes with DNA promoter sites and transcriptional regulators such as PhoB, which respond to phosphate starvation in a homologous manner across the entire bacterial kingdom.

By assembling complexes such as this, it is possible to interpret evolutionary conservation and covariation statistics mechanistically, providing important confirmation of limited mutation data.

We have also done this with eukaryotic protein kinases, and Oliver Lichtarge (Rice) has done this with GPCRs.



A structural model of the *E. coli* PhoB Dimer in the transcription initiation complex

Chang-Shung Tung\* and Benjamin H McMahon



# Large-scale biomolecular simulation of DNA with histones on Trinity supercomputer

Explicit-solvent simulation tests the bounds of large-scale molecular dynamics to investigate interplay of topology, entropy, and electrostatics to see how covalent modifications and DNA sequence-specific effects contribute to gene silencing.


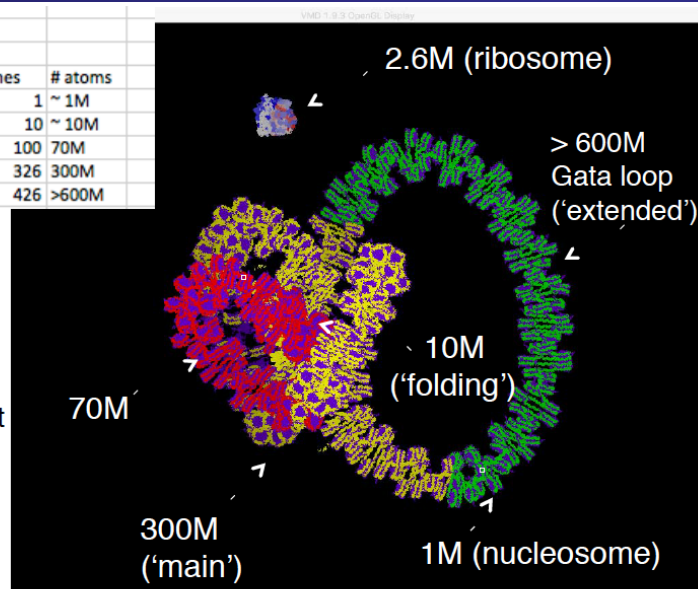
LANL: Karissa Sanbonmatsu, Mike Wall, Marcus Daniels, Toks Adedoyin, Wataru Nishima, Chang-Shung Tung

RIKEN: Jaewoon Jung, Chigusa Kobiyashi, Yuji Sugita


NYU: Gavin Bascom, Tamar Schlick

System Size			
System No.	System Name	# nucleosomes	# atoms
0	"benchmark"	1	~ 1M
1	"folding"	10	~ 10M
2	"Small" (red color)	100	70M
3	"Main" (yellow color)	326	300M
4	"Extended" (all colors)	426	>600M


- Largest run to date: 100M (Sugita 2016, milieu, GENESIS team)
- NYU visit (Spring 2016): GATA loop has high impact
- Explicit solvent runs stand to have high HPC impact
- Reduced model: more physics




# Pathway models consistent with physical chemistry have been extended via BioNetGen



## b Possible states of receptor subunits



## c A few of the 164 dimer states with active Syk



Molecular states are defined, accounting for both binding and phosphorylation

Parameters are (painfully) related to biochemical experiments, and the model is then propagated to the time-courses at right, where they can be fit to observations of system behavior.


Table I. Model parameter values<sup>a</sup>

Parameter	Value	Reference
<b>Components<sup>b</sup></b>		
Receptors, $R_T$	$4 \times 10^5/\text{cell}$	Current work
Lyn, $L_T$	$2.8 \times 10^4/\text{cell}$	Available Lyn $\sim 0.07 \times R_T$ (3)
Syk, $S_T$	$4 \times 10^5/\text{cell}$	Current work <sup>c</sup>
<b>Ligand binding</b>		
$k_{+1}$	$1.3 \times 10^{-10} \text{ molecules}^{-1} \text{ s}^{-1}$ ( $8 \times 10^4 \text{ M}^{-1} \text{ s}^{-1}$ )	Estimated binding parameters for covalently cross-linked IgE dimer (3)
$k_{+2}$	$2.5 \times 10^{-4} \text{ molecules}^{-1} \text{ s}^{-1}$	$k_{+2}R_T = 100 \text{ s}^{-1}$ (3)
$k_{-1}, k_{-2}$	0	$10^{-3} \text{ s}^{-1}$ over reported time scales
<b>Lyn association</b>		
$k_{+L}, k_{+L}$	$5 \times 10^{-5} \text{ molecules}^{-1} \text{ s}^{-1}$	Based on estimated equilibrium constants in Ref. 3
$k_{-L}$	$20 \text{ s}^{-1}$	Based on estimated equilibrium constants in Ref. 3
$k_{-L}^*$	$0.12 \text{ s}^{-1}$	Fit to observed rate of $\beta$ ITAM dephosphorylation from Ref. 51
<b>Syk association</b>		
$k_{+S}, k_{+S}$	$6 \times 10^{-5} \text{ molecules}^{-1} \text{ s}^{-1}$	Based on measured equilibrium constant at 25°C for binding of Syk tandem SH2 domains (46)
$k_{-S}, k_{-S}^*$	$0.13 \text{ s}^{-1}$	Fit to observed rate of $\gamma$ ITAM dephosphorylation in Ref. 51
<b>Phosphorylation</b>		
$p_{L\beta}, p_{L\beta}$	$30 \text{ s}^{-1}$	Consistent with extensive receptor phosphorylation
$p_{L\beta}^*, p_{L\beta}^*$	$100 \text{ s}^{-1}$	Moderate increase in Lyn kinase activity upon SH2 domain binding (15)
$p_{L\gamma}$	$1 \text{ s}^{-1}$	Double phosphorylation of $\gamma$ ITAM tyrosine, required to bind Syk, slower than single phosphorylation (48, 49, 77)
$p_{L\gamma}^*$	$3 \text{ s}^{-1}$	Moderate increase in Lyn kinase activity upon SH2 domain binding (15)
$p_{SS}$	$100 \text{ s}^{-1}$	Assumed same kinase activity as recruited Lyn
$p_{SS}^*$	$200 \text{ s}^{-1}$	Moderate increase in Syk kinase activity upon phosphorylation of the activation loop (24)
<b>Dephosphorylation</b>		
$d$	$20 \text{ s}^{-1}$	Fit to rates of ITAM dephosphorylation in Ref. 51

<sup>a</sup> Reactions associated with these rate constants are shown in Fig. 2.

<sup>b</sup> Cell density assumed to be  $1 \times 10^6$  cells/ml. Cell volume assumed to be  $1.4 \times 10^{-9}$  ml.

<sup>c</sup> Assayed value of  $S_T$  is  $3.4 \pm 0.4 \times 10^5/\text{cell}$ , which was rounded up so that  $S_T = R_T$ .




## Investigation of Early Events in FcεRI-Mediated Signaling Using a Detailed Mathematical Model<sup>1</sup>

James R. Faeder,\* William S. Hlavacek,\* Ilona Reischl,2\* Michael L. Blinov,\* Henry Metzger,‡ Antonio Redondo,† Carla Wofsy,\*§ and Byron Goldstein3\*

# Membrane protein simulations

In a multidisciplinary study of efflux pumps in *Pseudomonas* and *Burkholderia*, we developed an atomistic simulation capability of Gram negative membranes.

The interplay of influx, efflux, porins, quorum sensing, biofilm formation, and changes in gene expression was examined both theoretically and in laboratory experiments.



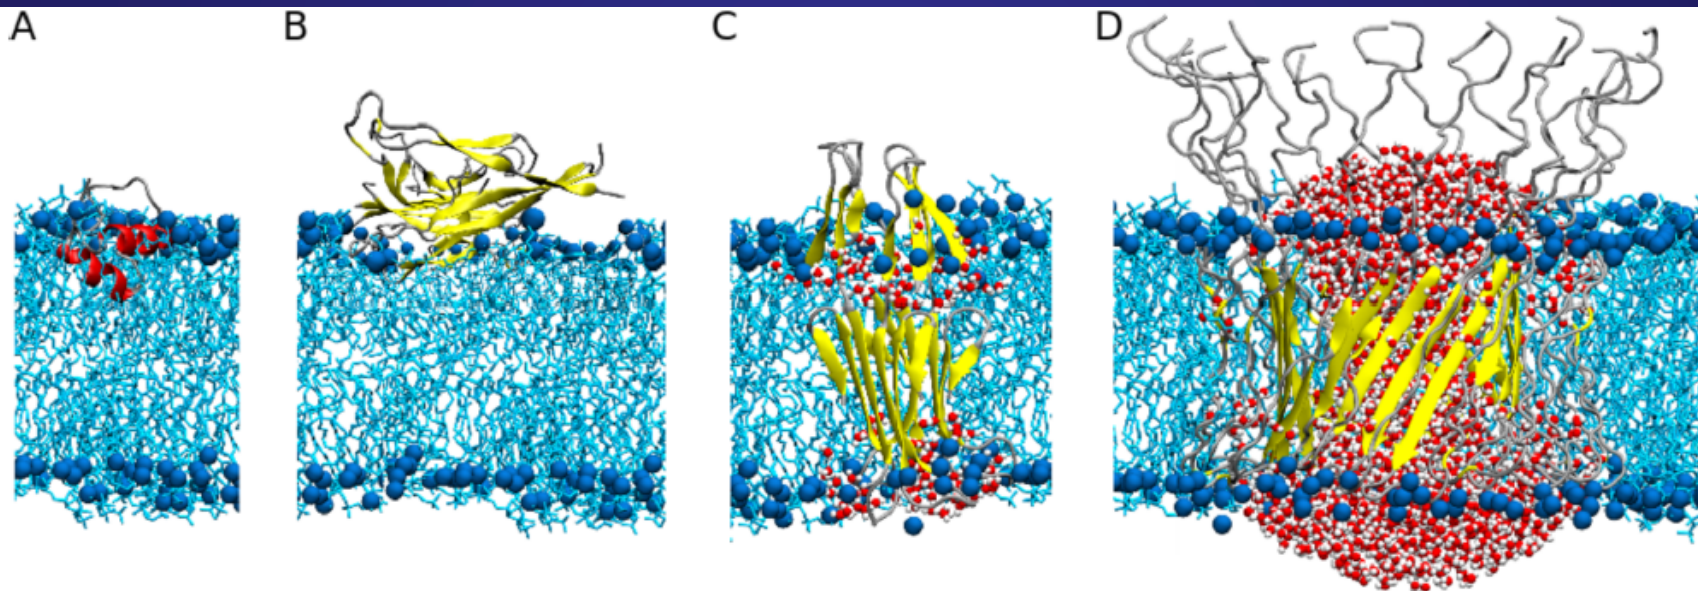
Permeability Barrier of Gram-Negative Cell Envelopes and Approaches To Bypass It

Helen I. Zgurskaya,<sup>\*,†</sup> Cesar A. López,<sup>§</sup> and S. Gnanakaran<sup>§</sup>




# Relating amyloid simulations to Alzheimer's disease

Relating molecular dynamics simulations to physiology typically requires extensive collaboration, such as described in this review article describing how amyloid simulations relate to Alzheimer's disease.




**Figure 9.** Various models of membrane-bound Aβ studied by MD simulations. (A) Monomeric Aβ1–40 is localized as a helical structure at the membrane surface. (B) Aβ1–40 forms a small cluster. (C) Aβ1–40 forms a larger cluster. (D) Aβ1–40 forms a large, dense cluster.



## Data analytics

- HIV databases and epidemiology
- HIV databases and vaccine research
- Mosaic vaccines
- DOE / NCI cancer genomics
- Signature-based metagenomics



# HIV and progress in epidemic surveillance

- Organized sampling
  - Two years after the establishment of the HIV sequence database, different subtypes were discovered
  - Viral recombination was characterized
  - It was recognized that culturing had a severe impact on the virus – only one phenotype could be cultured using standard cell lines
  - Several more to discover the multiple zoonotic events
- Quasispecies analysis and intervention
  - Initial drug resistance mutation surveillance was haphazard, many mutations were unrecognized, as was subtype dependence
  - Statistical technique to link mutations to phenotype changes existed but had never been applied to sequence data
  - Characterization of within-patient evolution and transmission took hundreds of publications
  - Recent work linking individual mutations to virus escape vs. lineage effects only possible for a few years
  - Ultra-rapid evolution and escape (order of weeks) only feasible with ultra-deep sequencing
  - Mapping of conserved and variable regions important for crystallography, vaccine design, drug design
- Surveillance
  - Epidemiological strain tracking was first applied in the early 90s
  - Pioneering work in attribution and network tracking
  - Initial geographical subtype distribution was mapped in the mid-90s (WHO initiative)
  - New subtypes and recombinants of epidemic relevance are named, published, and mapped
  - Recombinant composition often gives clues to origin of the strain
- New variants are now detected very quickly because of
  - Knowledge and vigilance in the field
  - Good, standardized assays at different sensitivity/specificity levels
  - Scientific interest and rewards
  - Dedicated infrastructure worldwide (partly due to drug trials & treatment programs)
  - Dozens of surveillance programs, collaborations and exchange programs between developed and developing world – human and simian subjects
  - Excellent, easy to use repository of existing strains also serves as knowledge and reference hub

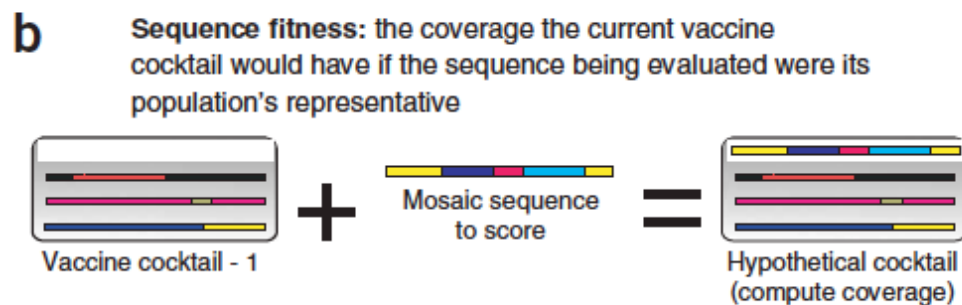
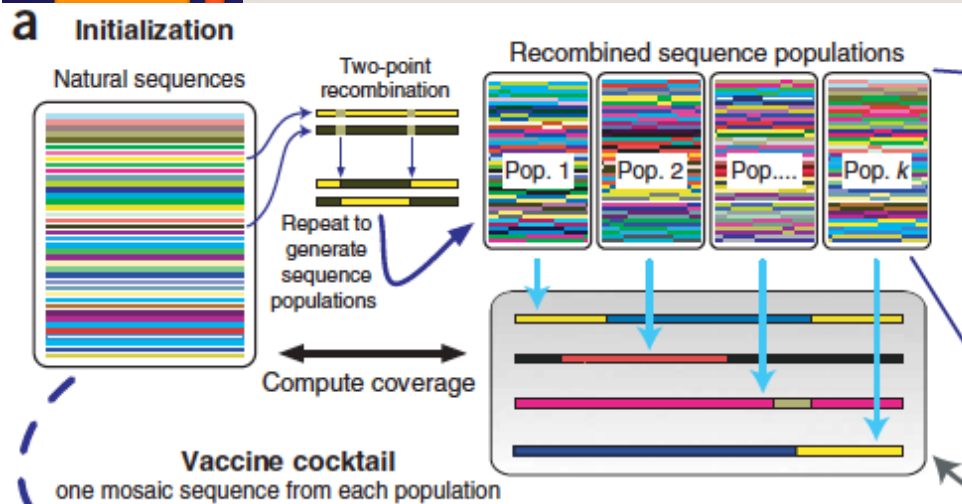


# How has this HIV knowledge helped?

- Distinguish evolution from recombination (attribution)
- Improve assays – detection, viral load – either more specific or more general
- Determine host cell specificity for individual virions
- Create virus-specific evolutionary model (epidemic age, zoonoses)
- Recognize new resistance mutations
- Avoid vaccine design mistakes
- Target vaccine to geography (could be done, though not a good strategy for HIV)
- Track sub-epidemics, determine origin of individual samples (prevent/mitigate panic)
- Monitor treatment efficacy, adapt regimens
- Standardize nomenclature and localization (e.g. mutations and epitopes)
- Researchers contribute to data repository out of self-interest



# Mosaic vaccine design



**Cocktail coverage:** the fraction of sequence 9-mers contained in the cocktail, averaged over all natural sequences

Polyvalent vaccines for optimal coverage of potential T-cell epitopes in global HIV-1 variants

Mosaic vaccines emerged from LANL's HIV database and associated analyses. They result in synthetic hybrid HIV proteins that maximally display commonly observed epitopes, and avoid confusing the immune system with rare epitopes.

The concept has been successfully extended to filoviruses, FMD, and HCV.

Leading to:




Figure 1. Protection of Rhesus Macaques Vaccinated with HIV-1 Mosaic Immunogens

David Palesch & Frank Kirchhoff

Will Fischer<sup>1,7</sup>, Simon Perkins<sup>1,7</sup>, James Theiler<sup>1</sup>, Tanmoy Bhattacharya<sup>1,2</sup>, Karina Yusim<sup>1</sup>, Robert Funkhouser<sup>1</sup>, Carla Kuiken<sup>1</sup>, Barton Haynes<sup>3</sup>, Norman L Letvin<sup>4</sup>, Bruce D Walker<sup>5</sup>, Beatrice H Hahn<sup>6</sup> & Bette T Korber<sup>1,2</sup>

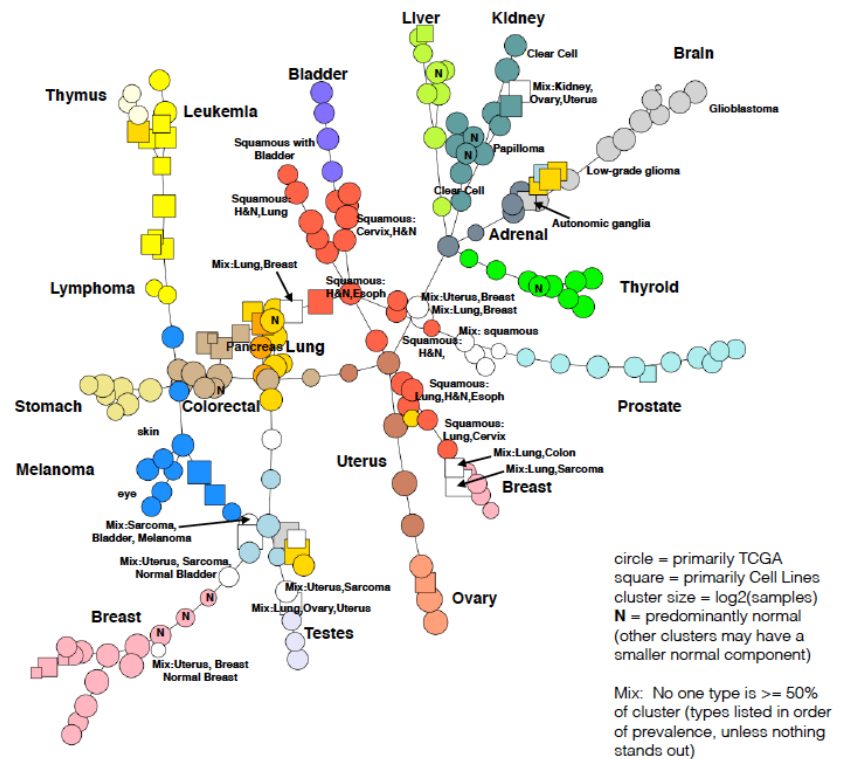


# DOE / NCI Cancer genomics

As part of the DOE/NCI precision medicine initiative, we are applying HPC to three problems in cancer medicine:

- Personalized genetics & mouse models
- Molecular understanding of RAS/RAF signaling
- Cancer epidemiology through the NCI/SEER program (registries)

The graph here provides insight into Q/C of cancer cell lines, as well as the most important distinctions of tissue type and cancer type. From here, alternative data types can be systematically integrated and properly analyzed through HPC.




Unsupervised clustering of 11,000 TCGA patient tumors and 2500 NCI, CCLE, and GDSC cancer cell lines on the basis of gene expression.

# Signature-based metagenomics


- $k$ -mers ( $N$ -grams, solid patterns) have a long history in both biological and natural-language research. In biology, they are physical objects.

English    Science is the great antidote to the poison of superstition.  
 French    La Science est le grand antidote au poison de la superstition.  
 Spanish    La ciencia es el gran antidoto al veneno de la superstición.  
 Dutch    Wetenschap is het grote tegengif aan het vergift van bijgeloof.  
 German    Wissenschaft ist das grosse Antidot zum Gift von Aberglauben.



## Amino-acid (protein) rather than nucleotide (DNA):

- Coding fraction is high for microbes
- No “silent” mutations-radius of convergence
- Functional pressures are mostly on proteins
- Protein sequences appear to be random strings
- 10-mers to achieve needed level of significance (< 99.7% non-random matches)
- Shared AA 10-mers are shared for a reason






# Decision Support

- Optimization of HCV combination therapies
- Pandemic simulations
- DoD Biosurveillance
- Patient trajectories

# Viral dynamics modeling to optimize HCV combination therapies.




Provides prediction of combination efficacy against HCV and rate of drug resistant strains appearing

**D**

1-nucleotide substitution

2-nucleotide substitution



Mechanistic infection model plus in-vitro assay

Quantifying antiviral activity optimizes drug combinations against hepatitis C virus infection

Yoshiki Koizumi<sup>a</sup>, Hirofumi Ohashi<sup>b,c</sup>, Syo Nakajima<sup>b,c</sup>, Yasuhito Tanaka<sup>d</sup>, Takaji Wakita<sup>b</sup>, Alan S. Perelson<sup>e</sup>, Shingo Iwami<sup>f,g,h,1,2</sup>, and Koichi Watashi<sup>b,c,h,1,2</sup>

# Models of targeted, layered, mitigations of a pandemic influenza run on a LANL supercomputer.


Table 2. Simulated mean number of ill people (cumulative incidence per 100) and for TAP, the number of antiviral courses required for various interventions and  $R_0$

Intervention	$R_0 = 1.6$	$R_0 = 1.9$	$R_0 = 2.1$	$R_0 = 2.4$
Baseline (no intervention)	32.6	43.5	48.5	53.7
Unlimited TAP (no. of courses)*	0.06 (2.8 M)	4.3 (182 M)	12.2 (418 M)	19.3 (530 M)
Dynamic vaccination (one-dose regimen) <sup>†‡</sup>	0.7	17.7	30.1	41.1
Dynamic child-first vaccination <sup>†‡</sup>	0.04	2.8	16.3	35.3
Dynamic vaccination (two-dose regimen) <sup>†§</sup>	3.2	33.8	41.1	48.5
Dynamic child-first vaccination <sup>†§</sup>	0.9	25.1	37.2	47.3
School closure <sup>¶</sup>	1.0	29.3	37.9	46.4
Local social distancing <sup>¶</sup>	25.1	39.2	44.6	50.3
Travel restrictions during entire simulation <sup>¶</sup>	32.8	44.0	48.9	54.1
Local social distancing and travel restrictions <sup>¶¶</sup>	19.6	39.3	44.7	50.5
TAP, * school closure, ** and social distancing <sup>**</sup>	0.02 (0.6 M)	0.07 (1.6 M)	0.14 (3.3 M)	2.8 <sup>††</sup> (20 M)
Dynamic vaccination, <sup>†‡</sup> social distancing, <sup>¶</sup> travel restrictions, <sup>¶¶</sup> and school closure <sup>**</sup>	0.04	0.2	0.6	4.5
TAP, * dynamic vaccination, <sup>†‡</sup> social distancing, <sup>¶</sup> travel restrictions, <sup>¶¶</sup> and school closure <sup>**</sup>	0.02 (0.3 M)	0.03 (0.7 M)	0.06 (1.4 M)	0.1 (3.0 M)
Dynamic child-first vaccination, <sup>†‡</sup> social distancing, <sup>¶</sup> travel restrictions, <sup>¶¶</sup> and school closure <sup>**</sup>	0.02	0.2	0.9	7.7

Physically meaningful parameters lead to specific, interpretable outcomes resulting from particular decisions. Uncertainty quantification is a major issue.

Updated model used to support DHS in responding to 2007 influenza pandemic, with 300,000,000 individuals belonging to households, schools, and workplaces, subject to rule-based mitigations and traveling according to USDoT statistics.

A similar framework can account for non-contagious disease progressions, with rule-based mitigations.

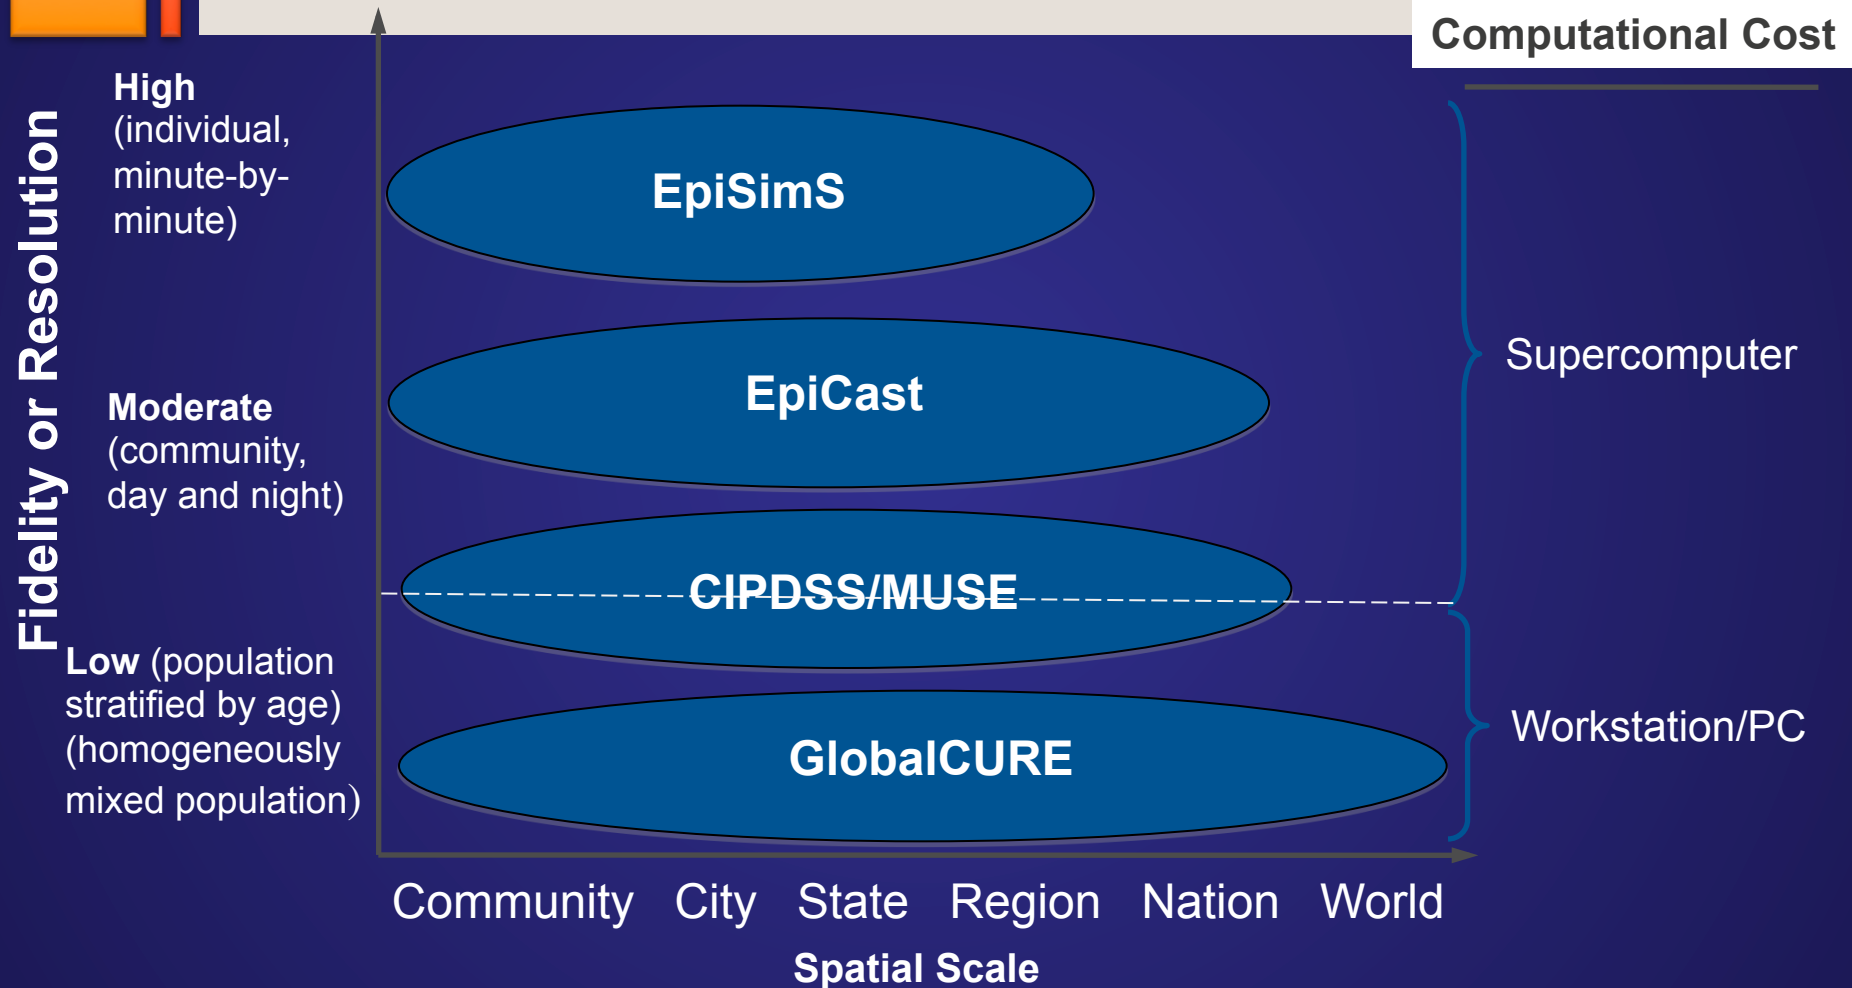


## Mitigation strategies for pandemic influenza in the United States

Timothy C. Germann<sup>\*,†</sup>, Kai Kadau<sup>\*</sup>, Ira M. Longini, Jr.<sup>‡</sup>, and Catherine A. Macken<sup>\*</sup>



# Epidemiological decision support for DoD




# VA patient timelines

Patient timelines can include multiple data-types:

- Vital signs (blood pressure, weight)
- Laboratory results (HbA1c, eGFR, LDL)
- Pharmaceuticals (insulin, oral agents)
- Hospital discharges, by ICD9/10 codes

At the VA, EMRs exist back to 2000, and are supplemented by numerous types of unstructured data, such as clinical notes, images, or discharge summaries.

The DOE / VA initiative will examine three medical areas – suicide prevention, prostate cancer outcome prediction, and cardio-vascular health, with a particular emphasis on genetic data from the Million Veteran Program.



An older effort attempts to identify diabetes treatment failure and apportion blame among patient attributes, provider attributes, and treatment attributes.



# Athena and HPC

Athena: Experimental multi-organ system  
for toxicology

High Performance Computing: Practical  
foundation rooted in multi-sponsor  
experience






# ATHENA



PI: Rashi Iyer, Ph.D.  
Los Alamos National Laboratory  
INTO-RAM EXEL Program




# Platform Engineering - Integration

## Flow Management




Liver flow controls  
at Charite



Athena Platform  
with Liver

- Developed pumps and valves for automation
- Developed Fluid Circuit boards – a multi-layer microchannel network to aid flow management
- Developed custom reservoirs for long term tissue culture, gas exchange, mixing, flow observation, etc.
- Developed software for control and automation
- Completed Athena platforms for Liver, Heart, and Lung BRs


## Organ specific developments




Video: Heart Assist

Pump on:  
Inspiration

Pump off:  
Expiration



Non-pneumatic actuation of MPS



Membrane  
development



Interconnected organs: Liver & heart



# Applied Genomics Capabilities

## LANL Computational Capabilities and Infrastructure

- **Capabilities**
  - Languages: R, Java, Perl, Python, C++, etc.
  - All mainstream open source software tools
  - A multitude of custom pipelines for analysis
- **Genome Science specific infrastructure**
  - >100 servers with >3000 cores, >12TB RAM
  - ~1.5 PB storage space
- **Notable resources**
  - Dell R910 (40 cores, 1TB RAM)
  - 2 x Dell 920 (80 cores, 1TB RAM)
  - IBM X3850 (80 cores, 2 TB RAM)
  - LANL HPC supercomputing



## Open Collaborative Network (Turquoise)

Name (Program <sup>1</sup> )	Processor	OS	Total Compute Nodes	CPU cores per Node / Total CPU cores	Memory per compute Node / Total Memory	Interconnect	Peak (TFlop/s)	Storage
Grizzly CTS-1 <sup>2</sup> (IC)	Intel Xeon Broadwell	Linux (TOSS)	1490 nodes	36 / 53,640	128 GB / 191 TB	Intel OmniPath	1,806	15.2 PB Lustre
Lightshow TLCC <sup>3</sup> (ASC)	Intel Xeon X5650 and 2x nVidia Quadro GPUs	Linux (TOSS)	16 nodes	12 / 192	96 GB / 1.5 TB	Mellanox Infiniband Fat-Tree	4.0	15.2 PB Lustre
Moonlight TLCC2 <sup>3</sup> (ASC)	Intel Xeon E5-2670 + 2x NVidia Tesla M2090 GPUs	Linux (TOSS)	308 nodes	16 / 4,928 + GPUs	32 GB / 9.86 TB	Qlogic InfiniBand Fat-Tree	488	15.2 PB Lustre
Pinto TLCC2 <sup>3</sup> (IC)	Intel Xeon E5-2670	Linux (TOSS)	154 nodes	16 / 2,464	32 GB / 4.9 TB	Qlogic InfiniBand Fat-Tree	51.3	15.2 PB Lustre
Wolf TLCC2 <sup>3</sup> (IC)	Intel Xeon E5-2670	Linux (TOSS)	616 nodes	16 / 9,856	64 GB / 39.4 TB	Qlogic InfiniBand Fat-Tree	205	15.2 PB Lustre

<sup>1</sup> Programs: IC=Institutional Computing, ASC=Advanced Simulation and Computing, R=Recharge <sup>2</sup> CTS-1 = Commodity Technology System, 1st Generation. <sup>3</sup> TLCC = TriLab Linux Capacity Cluster; 2 = 2nd Generation <sup>4</sup> ATS-1 = Advanced Technology System, Generation 1

## Secret Restricted Network (SRD - Red)

Name (Program <sup>1</sup> )	Processor	OS	Total Compute Nodes	CPU cores per Node / Total CPU cores	Memory per compute Node / Total Memory	Interconnect	Peak (TFlop/s)	Storage
<b>Fire</b> CTS-1 <sup>2</sup> (ASC)	Intel Xeon Broadwell	Linux (TOSS)	1104 nodes	36 / 39,744	128 GB / 141 TB	Intel OmniPath	1,338	9.3 PB Lustre
<b>Ice</b> CTS-1 <sup>2</sup> (ASC)	Intel Xeon Broadwell	Linux (TOSS)	1104 nodes	36 / 39,744	128 GB / 141 TB	Intel OmniPath	1,338	9.3 PB Lustre
<b>Luna</b> TLCC2 <sup>3</sup> (ASC)	Intel Xeon Sandybridge	Linux (TOSS)	1,540 nodes	16 / 24,640	32 GB / 49 TB	Qlogic InfiniBand Fat-Tree	513	9.3 PB Lustre
<b>Trinity</b> ATS-1 <sup>5</sup> (ASC)	Intel Xeon Haswell + Knights Landing	SLES-based CLE and CNL	9,408 Xeon Haswell + 9500 KNL nodes	32 XH + 68 KNL / 947,056	128GB XH + 96GB KNL / 2.1 PB	Cray Aries Dragonfly	42,170	72 PB Lustre

# Unclassified Protected Network (Yellow)

Name (Program <sup>1</sup> )	Processor	OS	Total Compute Nodes	CPU cores per Node / Total CPU cores	Memory per compute Node / Total Memory	Interconnect	Peak (TFlop/s)	Storage
Snow CTS-1 <sup>2</sup> (ASC)	Intel Xeon Broadwell	Linux (TOSS)	368 nodes	36 / 13,248	128 GB / 47.1 TB	Intel OmniPath	445	15.2 PB Lustre
Trinitite (ASC)	Intel Haswell (HSW) and Knights Landing (KNL)	SLES-based CLE and CNL	100 HSW nodes + 100 KNL nodes	32 / 3,200 HSW + 68 / 6,800 KNL	HSW 128 GB / 12.8 TB + KNL 112 GB / 11.2 TB	3D Aries Dragonfly	364	678 TB Lustre





Delivering science and technology  
to protect our nation  
and promote world stability